HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P. O. Box 272400
Fort Collins, Colorado 80527-2400

PATENT APPLICATION IFW

ATTORNEY DOCKET NO. 10992806-1

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

AP

Inventor(s): Robert Alan Cochran

Confirmation No.:

Application No.: 09/678,168

Examiner: Akiba K.R. Boyce

Filing Date: 10/02/2000

Group Art Unit: 3623

Title: Method and System for Throttling I/O Request Servicing on Behalf of an I/O Request Generator to Prevent Monopolization of a Storage Device by the I/O Request Generator

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

12/07/2004 EFLORES 00000014 09678168
01 FC:1251 110.00 DA

TRANSMITTAL OF APPEAL BRIEF

Sir:

Transmitted herewith in ***triplicate*** is the Appeal Brief in this application with respect to the Notice of Appeal filed on ***Aug. 30, 2004***.

The fee for filing this Appeal Brief is (37 CFR 1.17(c)) $330.00.

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

(X) (a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d) for the total number of months checked below:

(X)	one month	$110.00
( )	two months	$420.00
( )	three months	$950.00
( )	four months	$1480.00

( ) The extension fee has already been filled in this application.

( ) (b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account ***08-2025*** the sum of ***$440.00***. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

( ) I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, Alexandria, VA 22313-1450. Date of Deposit: ***Nov. 30, 2004***

OR

( ) I hereby certify that this paper is being transmitted to the Patent and Trademark Office facsimile number_____ on _____

Number of pages:

Typed Name: ***Joanne Bourguignon***

Signature: _____

Respectfully submitted,

Robert Alan Cochran

By _____

Robert W. Bergstrom

Attorney/Agent for Applicant(s)
Reg. No. ***39,906***

Date: ***Nov. 30, 2004***

Telephone No.: ***206.621.1933***

12/07/2004 EFLORES 00000014 082025 09678168
01 FC:1251 110.00 DA

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re patent application of:

Inventors: Robert Alan Cochran
Serial No. 09/678,168
Filed: October 02, 2000
For: Method and System for Throttling I/O Request Servicing on Behalf of an I/O Request Generator to Prevent Monopolization of a Storage Device by the I/O Request Generator

Examiner: Akiba K.R. Boyce
Group Art Unit: 3623
Docket No. 10992806-1
Date: November 30, 2004

APPEAL BRIEF

Commissioner of Patents and Trademarks
Washington, DC 20231

Sir:

This appeal is from the decision of the Examiner, in an Office Action mailed on April 05, 2004, finally rejecting claims 1-20.

REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

RELATED APPEALS AND INTERFERENCES

Applicant's representative has not identified, and does not know of, any other appeals of interferences which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

STATUS OF CLAIMS

Claims 1-20 are pending in the application. Claims 1-20 were finally rejected in the Office Action dated April, 05, 2004. Applicant's appeal the final rejection of claims 1-20, which are copied in the attached Appendix I.

STATUS OF AMENDMENTS

No Amendment After Final is enclosed with this brief. The last Response was filed June 7, 2004. The last amendment to the claims was filed December 19, 2003.

SUMMARY OF CLAIMED SUBJECT MATTER

Overview

In one embodiment, the present invention is employed by a disk-array controller to prevent monopolization, by remote computers that elect a premium tier of servicing, of I/O request servicing provided by the disk-array controller to remote computers. In other words, the disk-array controller needs to provide the expected premium-tier servicing of I/O requests from premium-tier remote computers, but needs also to make sure that non-premium-tier computers receive a fair amount of I/O request servicing.

Independent Claim 1

Claim 1 is directed to a method employed by a request servicing device for fairly servicing electronic requests received from request generating devices interconnected with the request receiving device. A pricing tier is established for each request generating device. A maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request servicing, is established for the request servicing device. The request servicing device maintains an instantaneous rate of request servicing by the request servicing device for each request generating device with a premium pricing tier. After the

request servicing device services a request of a request generating device, the request servicing device determines a time elapsed during servicing of the request. When the determined time elapsed during servicing of the request is less than the expected time for servicing the request, the request servicing device calculates a remaining time equal to the difference between the expected time and the elapsed time and waits for a length of time based on the calculated remaining time prior to servicing another request for the request generating device.

Dependent Claims 2 – 9

Claim 2 provides details of how the time to wait before servicing another request from a particular request generating device is computed. Claim 3 is directed to the method of claim 1 wherein the request generating device is a computer. Claim 4 is directed to the method of claim 1 wherein the request servicing device is a data-storage device. Claim 5 is directed to the method of claim 1 wherein the request servicing device is a disk array data-storage device. Claim 6 is directed to the method of claim 1 wherein a request generating device specifies the maximum rate of request servicing that it desires. Claim 7 is directed to the method of claim 1 wherein the request servicing device partitions its request servicing bandwidth among the request generating devices, and employs this partitioning to establish a maximum rate of request servicing for each request generating device. Claim 8 is directed to the method of claim 1 wherein the request servicing device dynamically alters the maximum rate of request servicing for a request generating device. Claim 9 is directed to the method of claim 1 wherein the request servicing device dynamically alters the instantaneous rate of request servicing for a request generating device after servicing a request in an efficient manner.

Independent Claim 1

Claim 10 is directed to a request servicing device that fairly service electronic requests received from remote request generating devices. The request servicing device includes a memory that store various parameters associated with each remote request generating device, including a pricing tier, a maximum rate of request servicing, and an expected time for servicing a request at the maximum rate of request servicing. For premium-tier remote request generating devices, the request servicing device continuously computes and maintains in the memory an instantaneous rate of request servicing. The

request servicing device, upon servicing each request, determines the time taken to service the request. When this time is less than the expected time for servicing the request, the request servicing device computes a wait time based on the elapsed time and waits for that wait time before servicing another request for the remote request generating device.

Dependent Claims 11 – 20

Claim 11 provides details of how the time to wait before servicing another request from a particular request generating device is computed by the request servicing device. Claim 12 is directed to the method of claim 1 wherein the request generating device is a computer. Claim 13 is directed to a data-storage-device request servicing device. Claim 14 is directed to a disk array request servicing device. Claim 15 is directed to a request servicing device that allows a request generating device to specific the maximum rate of request servicing that it desires. Claim 16 is directed to a request servicing device that partitions its request servicing bandwidth among the request generating devices, and employs this partitioning to establish a maximum rate of request servicing for each request generating device. Claim 17 is directed to a request servicing device that dynamically alters the maximum rate of request servicing for a request generating device. Claims 18-19 are directed to a request servicing device that dynamically alters the instantaneous rate of request servicing for a request generating device after servicing a request in an efficient manner. Claim 20 is directed to a request servicing device that computes the length of time to wait before servicing another request for a particular request generating device by an efficient method.

GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1, 6, 10, 15, and 20 stand rejected under 35 U.S.C. § 102(e) as being anticipated Kilkki et al., U.S. Patent No. 6,011,778 ("Kilkki").

2. Claims 2-3 and 11-12 stand rejected over Kilkki under 35 U.S.C. § 103(a).

3. Claims 4-5, 7-9, and 13-19 stand rejected over Kilkki in further view of Storch et al., U.S. Patent No. 5,920,846 ("Storch") under 35 U.S.C. § 103(a).

ARGUMENT

ISSUE 1

1. Whether claims 1, 6, 10, 15, and 20 are anticipated, under 35 U.S.C. § 102(e), by Kilkki et al., U.S. Patent No. 6,011,778 ("Kilkki").

Claim 1 is representative of the group of claims including claims 1, 6, 19, 15, and 20. Claim 1 is provided below, with added emphasis:

> 1. A method for fairly servicing, by a request servicing device, electronic requests received by the request servicing device from request generating devices interconnected with the request receiving device, the method comprising:
>
> establishing a pricing tier for each request generating device, a maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request servicing;
>
> for each request generating device with a premium pricing tier, maintaining an instantaneous rate of request servicing by the request servicing device;
>
> following servicing of each request from a request generating device by the request servicing device,
>
> determining a time elapsed during servicing of the request;
>
> when the time elapsed during servicing of the request is less than the expected time for serving a request established for the request generating device,
>
> calculating a remaining time equal to the difference between expected time for serving a request established for the request generating device and the time elapsed during servicing of the request; and
>
> waiting for a length of time based on the calculated remaining time prior to servicing another request for the request generating device.

Claim 1 is directed to a method for servicing electronic requests employed by a request servicing device. In an embodiment described in the current application in detail, the request-servicing device is a disk-array controller running within a disk array. The disk-array controller receives I/O requests from remote computers and services the received requests, returning data and/or status related to the I/O requests to the requesting, remote computers. As clearly claimed in claim 1, the request servicing device, or disk-array controller in the described embodiment, establishes and maintains, for each remote computer on behalf of which the request servicing device services I/O requests, a pricing tier, a maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request servicing. After the request servicing device services a request on behalf of a remote computer, the disk-array controller computes an elapsed time for servicing of the request.

When the elapsed time is less than the expected time for servicing the I/O request, the request servicing device computes a difference between the elapsed time and the expected time and waits for a time equal to the computed difference before servicing another I/O request on behalf of the remote computer.

Kilkki, by contrast, is directed to a timer-based method employed by an access node for measuring transmission rates of information by a network and for adjusting message priority assignments by the access node to achieve a nominal bit rate for transmission by the network. Kilkki's method is best understood with reference to Kilkki's Figure 2. Kilkki's disclosed system involves a user (Figure 2, 20) that transmits messages, or, in Kilkki's terminology, transmits cells, through a user/network interface ("UNI") (Figure 2, 24) to a first node (Figure 2, 32) in a computer network (Kilkki, column 7, lines 16-18) that, in turn, transmits the cells to additional computers (Figure 2, 34) in the network; a final network node 34 forwarding the cells to destination computers (Figure 2, 36). The user interacts, through the UNI 24, with a network operator (Figure 2, 22), presumably another computer, to negotiate (or, presumably, to purchase) a nominal bit rate ("NBR") for transmitting data through the network (Kilkki, column 7, lines 45-46). Kilkki does not appear to explicitly state whether the UNI is a software program running on the user's computer, or is instead a separate access node distinct from the user's computer.

The Examiner seems to indicate, in the Advisory Action dated July 28, 2004, that the Examiner appears to consider the user and UNI together as a request generating device for purposes of comparing Kilkki to the currently claimed invention, stating:

> As stated in the rejection, the NBR (nominal bit rate) connection by way of UNI (user network interface as shown in Col. 5, lines 60-63) over an NBR service connection, represents request generating device since the user is shown to make an NBR request via the network in col. 6, lines 36-42.

This is not the only possible mapping of computers in Kilkki's system to claimed components of the current invention, but is a reasonable mapping. Other mappings might consider the user to be the request generator, and the UNI to be the request servicing device. Defining the user and UNI together as the request generating device, as the Examiner appears to have done, does not produce an interpretation of Kilkki's system on which claim 1, or any other claim in the current application, reads, and instead produces a system quite unrelated to the currently claimed invention, as discussed below. Claim 1 does not read on any other mapping that Applicant's representative can imagine.

Taking the user/UNI combination as the request generating device, one needs to next determine what, in Kilkki's system, constitutes a request servicing device. There does not appear to be a single request servicing device. The network operator (Figure 2, 22) services requests for assignment or purchase of a nominal bit rate, but does not service requests for transmission of cells to destination computers. The first network node 32 directly receives cell-transmission requests from the user/UNI request generating device. However, the first network node 32 does not, by itself, service the cell-transmission requests. Instead, the first network node transmits a cell to the network, and additional network nodes forward the cell from one to another, finally routing the cell to a destination computer. The destination computer may also be a request servicing device, if the cells transmitted to the destination computer encode higher-level requests of some kind. There is, however, no mention of request servicing by the destination computer in Kilkki. It is probably most reasonable to map the claimed request servicing device of claim 1 to the first network node (Figure 2, 32), since the request generating device, or user/UNI, is disclosed in Kilkki as generating cell-transmission requests and forwarding those requests to the first network node.

In Kilkki's system, the UNI attempts to achieve the nominal bit rate ("NBR"), negotiated for with the network operator by the user, for transmitting data through the network. The NBR represents an expected, but not a guaranteed, bit rate associated with the user (Kilkki, column 3, lines 32 – 34). The UNI 24 sets a timer to an expected time upon transmitting a first cell to the first network node 32. The timer value is evaluated by the UNI at the point in time that the UNI transmits a second cell to the first network node. The evaluation of the timer value provides a difference between the expected time interval between transmission of the first and second cells and the actual time interval between transmission of the first and second cells. In other words, the elapsed time calculated *by the request generating device* is an elapsed time between issuing electronic cell-transmission requests by the *request generating device.* The UNI assigns a priority level to the second message based on this difference (Kilkki, column 3, lines 53-65; column 6, lines 11-35). The first network node 32 receives the cell from the UNI, and decides, based on the priority assigned to the cell by the UNI, whether or not to transmit the cell through the network, or to instead discard the cell:

> After computing the priority level of each cell at the UNI, the cells are transmitted 48 to the network, such as to a node of the network. A network node, upon arrival of a cell transmitted from the UNI, performs a cell filtering process by which the node determines whether to accept or discard a particular cell. (Kilkki,

column 6, lines 21-26)

Claim 1, provided above, claims a method for fairly servicing, *by a request servicing device*, electronic requests received by the request servicing device from request generating devices interconnected with the request receiving device. In the first step of the claimed method, the request servicing device (1) establishes a pricing tier for each request generating device; (2) establishes a maximum rate of request servicing; and (3) establishes an expected time for serving a request at the maximum rate of request servicing. By contrast, in Kilkki's disclosed system, the network operator (Figure 2, 22), and not the request servicing device (first network node 32), establishes the NBR, or, in other words, the maximum rate of request servicing. In Kilkki's system, the *request generating device* (user 20 and UNI 24) computes an *expected time between issuing cell-transmission requests* – the only expected time mentioned in Kilkki. Note that this is not an expected time for servicing a particular request, or for servicing requests in general, as in the currently claimed method,

It is already clear that Kilkki's disclosed method and system cannot possibly anticipate the method of claim 1. According to MPEP § 2131:

> "A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." ... "The identical invention must be shown in as complete detail as is contained in the ... claim."

In particular, the request servicing device, in the claimed method, establishes a pricing tier, a maximum rate of request servicing, and an expected time for servicing a request. In Kilkki's system, by contrast, neither a request generating device nor a request servicing device establishes the NBR. In Kilkki's system, the request generating device, rather than the request servicing device, establishes an expected time. However, the expected time is an expected time interval between making cell-transmission requests of the first network node by the user/UNI, and is not an expected time for servicing a request.

It should be noted that Kilkki explicitly states, in several places, that the network node, or request servicing device, needs to know very little about the user/UNI devices:

> The SIMA service concept, from a user's perspective, is simple and understandable, because there is no pre-defined traffic or quality parameters associated with each connection, and charging for connection usage is based solely on the value of the NBR and the duration of the connection.
>
> A typical implementation of a SIMA service utilizes two primary components: access nodes and core network nodes, which have fundamentally different functional

> responsibilities. For example, access nodes, which may be a user/network interface, perform the task of measuring traffic for every connection, whereas at the core network nodes, the traffic control functions do not need to know anything about the properties of individual connections. (Kilkki, column 5, lines 11-24)
>
> Although a user may change any of these parameters, the only information that the network needs to know at the initiation of cell transmission is the NBR and the service class (real-time or non-real-time) of the connection. (Kilkki, column 11, lines 29-33)

In other words, the first network node 32 knows almost nothing about user/UNIs. In particular, the request servicing device in Kilkki does not compute or maintain information with respect to expected service times or service rates. As discussed above, any such activities are the responsibility of the *request generating device* in Kilkki's system.

The second step of claim 1, provided above, states that the *request servicing device* maintains an instantaneous rate of request servicing by the request servicing device for each request generating device with a premium pricing tier. In the Office Action, the Examiner referred to Kilkki, column 11, lines 22-24 for teaching maintenance of an instantaneous rate of request servicing by the request servicing device. However, these cited lines state that: "In addition, it may be desirable to request that each user select an appropriate instantaneous NBR, which should be no greater than the selected maximum NBR." The second step of claim 1 refers to the *request servicing device*, or disk-array controller in the embodiment described in the current application, maintaining an instantaneous rate of request servicing by the *request servicing device*. In the described embodiment, the request servicing device computes and maintains, based on a sliding-window history of request servicing, a recent request servicing rate with respect to each remote client. That is quite different from requesting that a user supply a desired instantaneous NBR to the network. As discussed in the preceding paragraph, the request servicing device in Kilkki's system maintains almost no information about remote user/UNIs. There is no indication that even the request generating device in Kilkki's system, a user/UNI, computes and maintains an instantaneous rate of request servicing in Kilkki's system. At most, in Kilkki's system, the user/UNI computes a difference between an interval from sending a first request and sending a second request and an expected interval between sending the first and second requests.

In the final steps of claim 1, when the time elapsed during servicing of a request is less than the expected time for serving a request established for the request generating device *by the request servicing device*, the request servicing device, or disk-array

controller in the described embodiment, calculates a remaining time equal to the difference between the expected time for serving a request established for the request generating device and the time elapsed during servicing of the request, and *waits for a length of time based on the calculated remaining time* prior to servicing another request for the request generating device. In Kilkki's system, neither the request generating device (user/UNI) nor the request servicing device (network node) waits for any length of time. In Kilkki's system, the request generating device computes a priority to associate with the next cell-transmission request based on the time interval between sending that request and sending the previous request. The request generating device then sends the priority-associated cell-transmission request. It does not wait for a computed period of time before doing so. In Kilkki's system, the network node determines, based on the priority associated with the cell-transmission request, whether to transmit the cell or discard the cell, as discussed in the above quoted passage from Kilkki (Kilkki, column 6, lines 21-26). The request servicing device does not wait for a computed length of time. The request servicing device in Kilkki may, based on the priority assigned to a cell, not service the request at all.

Applicant's representative can find no way to read claim 1 onto Kilkki's disclosed method and system. Claim 10 includes language similar to the language of claim 1, and is not anticipated by Kilkki for the reasons provided above with respect to claim 1. It is not surprising that Kilkki does not anticipate claims 1 and 10, and the dependent claims that depend from them, since Kilkki discloses a method and system entirely different from the claimed invention and embodiments of the claimed invention disclosed in the current application. The claimed invention involves a method, and a system embodying the method, by which a request servicing device prevents monopolization of request servicing bandwidth provided by the request servicing device by a remote request generating device, as clearly stated even in the title of the current application. By contrast, Kilkki discloses a method by which a request generating device attempts to ensure that the request generating device obtains adequate request servicing from a request-servicing, multiple-device system. Kilkki's method is directed to enabling a request generating device to obtain a level of request servicing that a user has contracted for, while the current invention, clearly claimed in claims 1 and 10, and clearly described in the current application, is directed to enabling a request servicing device to prevent monopolization of the request servicing device by premium-tier request generating devices while attempting to meet request servicing rates appropriate for each request generating device.

ISSUE 2

2. Whether claims 2-3 and 11-12 are unpatentable, under 35 U.S.C. § 103(a), over Kilkki.

As discussed above, with respect to Issue 1, the claimed invention involves a method, and a system embodying the method, by which a request servicing device prevents monopolization of the request servicing bandwidth, provided by the request servicing device, by a remote request generating device. By contrast, Kilkki discloses a method in which a request generating device attempts to ensure that the request generating device issues requests, to a request-servicing, multiple-device system, at a rate equivalent to a negotiated nominal bit rate by assigning priorities to requests issued to the request-servicing, multiple-device system.

The Examiner apparently feels that Kilkki discloses calculation of a remaining time for a request generating device in a way different from something included in claims 2-3 and 11-12, and therefore, in Section 6 of the Office Action, rejected these claims based on 35 U.S.C. § 103(a), rather than on 35 U.S.C. § 102(e). According to MPEP § 2143:

> To establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations.

Applicant's representative cannot parse, and does not understand, the Examiner's stated rationale for this rejection. However, there are far more differences between the claimed invention and the method and system disclosed by Kilkki than any single difference identified by the Examiner. Kilkki's disclosed system and method are far too different from the claimed invention to provide a basis for an obviousness-type rejection.

As one example of the great difference between Kilkki's disclosed system and method of the claimed invention, claim 2 provides that the request servicing device computes the time to wait before servicing another request on behalf of a premium-tier request generating device, when the computed rate of request servicing is below the maximum rate of request servicing, to be "less than the calculated remaining time for a request generating device." In other words, the request servicing device may boost the instantaneous rate of

request servicing above the contracted maximum request servicing rate in order to provide an overall request servicing rate close to the maximum request servicing rate. This is described on lines 16-27 of the current application. In Kilkki's system and method, as discussed with respect to Issue 1, the request servicing device does not maintain such information, and does not control request servicing rates. As discussed above, in Kilkki's system, it is the *request generating device* that is responsible for attempting to ensure an NBR. Moreover, in the only mention of an instantaneous request servicing rate, Kilkki states: "In addition, it may be desirable to request that each user select an appropriate instantaneous NBR, which should be no greater than the selected maximum NBR" (Kilkki, column 11, lines 21-24). Kilkki's disclosed method and system are at best unrelated to the current invention, and, with respect to instantaneous rates of request servicing, actually teach away from the claimed invention. Kilkki's user-selected instantaneous NBR should not exceed the contracted for, maximum NBR.

ISSUE 3

3. Whether claims 4-5, 7-9, and 13-19 are unpatentable, under 35 U.S.C. 103(a), over Kilkki in further view of Storch et al., U.S. Patent No. 5,920,846 ("Storch").

As discussed above, with respect to Issue 1, the claimed invention involves a method, and a system embodying the method, by which a request servicing device prevents monopolization of the request servicing bandwidth, provided by the request servicing device, by a remote request generating device. By contrast, Kilkki discloses a method in which a request generating device attempts to ensure that the request generating device issues requests, to a request-servicing, multiple-device system, at a rate equivalent to a negotiated nominal bit rate by assigning priorities to requests issued to the request-servicing, multiple-device system. Kilkki's disclosed system and method are quite dissimilar in approach, implementation, and goals, from Applicant's clearly disclosed and claimed invention. Kilkki, by itself, neither anticipates nor makes obvious the claimed invention.

Storch is entirely unrelated art, as can be easily determined from a cursory reading of the initial lines of Storch's Abstract:

> An integrated system and method is provided for processing a service request for installation, maintenance or repair of a local loop maintained by a telecommunications company and providing locally switched service to a customer premise. The system includes a computer network that allows for inputting and

> processing customer information; generating a service request based upon customer information and information associated with the local loop; assigning to the service request an available appointment time for providing the requested service based upon updated information indicating the availability of qualified outside technicians; dispatching an available technician to the customer premise at or before the assigned appointment time to install or repair outside facilities associated with the local loop or to install or replace a special line conditioning termination equipment at a network interface for the customer premise when the loop loss associated with the local loop is not maintained at an acceptable level for the locally switched service, processing completion information input by the technician ...

As is readily apparent, Storch's method and system is related to servicing telephone customers by sending out repair technicians. In Storch, the request servicing entity is a repair technician. Applicant's representative has, in good faith tried to find any kind of similarity between the subject matter disclosed by Storch and the current invention, but can find none.

As one example of the inapplicability of Storch, the Examiner states, in Section 7 of the Office Action:

> As per claim 7, 16, Kilkki et al fails to disclose, however Storch et al discloses:
>
> wherein the maximum rate of request servicing established by partitioning the capacity of the request servicing device among the request generating devices in order to provide, when possible, each request generating device with a maximum rate of request servicing specified by the request generating device, and otherwise to provide each request generating device with a maximum rate of request servicing proportional to a maximum rate of request servicing specified by the request generating device, (Col. 58, line 59-Col. 60, line 2, [line intervals being divided]).

The cited lines of Storch begin with the sentences: "Based upon the information transmitted by the DUDAS 266, SORD 268 data processes this information to continually update records stored in its computer memory indicating whether appointments are open, closed, unavailable (not offered), or the time falls on a holiday, and transits this information to the computer order entry system 254 to assist the order taker person 252 in scheduling appointments. The appointment times can be defined as any time interval (such as A.M. and P.M. appointments where the A.M. and P.M. time interval can extend between any predefined time limits, or divided into two hour increments)." The cited lines of Storch end with the sentence: "These service codes are industry standards used nationwide that provide a universal definition indicating the service provided to the customer." Applicant's representative is completely perplexed as to how these lines of Storch are, in any way, related to disk arrays or other such electronic request serving devices, and how the cited passage from Storch relates to the

language referenced by the Examiner in the Office Action that includes phrases like "maximum rate of request servicing by a request servicing device."

> Also in Section 7 of the Office Action, the Examiner states:
>
> As per claims 5, 14, both Kilkki et al and Storch et al fail to disclose: wherein the electronic storage device is a disk array.
>
> Official notice is taken that it is old and well known in the computer art for an electronic data storage device to be a disk array. It would have been obvious to one of ordinary skill in the art at the time of the applicant's invention for the electronic data storage device to be a disk array with the motivation of having reliable and easily accessible means for retrieving stored data.

Applicant's representative is again completely perplexed by this statement. Kilkki discloses a network communications system, and does not mention electronic data storage devices per se. Storch is completely unrelated to the current application, involving telephone repair scheduling by repair technicians. Claim 5 depends from claim 4, further qualifying the phrase "request servicing device." Claim 5 essentially claims the method in which the request servicing device is a disk array. In Kilkki, the request servicing device is a computer network node – not a disk array. In Storch, the request servicing entity is a repair technician – not a disk array. The above statement appears to be wholly without context or meaning with respect to the cited references.

Kilkki neither anticipates nor makes obvious the claimed invention, for reasons detailed above in the discussion of Issue 1. Storch is quite unrelated art, and adds nothing in combination with Kilkki. Therefore, neither independent nor dependent claims of the current application are made obvious by the combination of Kilkki and Storch.

CONCLUSION

The claimed invention involves a method, and a system embodying the method, by which a request servicing device prevents monopolization of the request servicing bandwidth, provided by the request servicing device, by a remote request generating device. By contrast, Kilkki discloses a method in which a request generating device attempts to ensure that the request generating device issues requests, to a request-servicing, multiple-device system, at a rate equivalent to a negotiated nominal bit rate by assigning priorities to requests issued to the request-servicing, multiple-device system. Storch is completely unrelated art. Neither Kilkki, Storch, nor Kilkki and Storch in combination, teach, mention, or suggest the claimed invention.

Applicant respectfully submits that all statutory requirements are met and that the present application is allowable over all the references of record. Therefore, Applicant respectfully requests that the present application be passed to issue.

Respectfully submitted,
Robert Alan Cochran
OLYMPIC PATENT WORKS PLLC

By _____
Robert W. Bergstrom
Reg. No. 39,906

Olympic Patent Works PLLC
P.O. Box 4277
Seattle, WA 98104
206.621.1933 telephone
206.621.5302 fax

CLAIMS APPENDIX

1. (original) A method for fairly servicing, by a request servicing device, electronic requests received by the request servicing device from request generating devices interconnected with the request receiving device, the method comprising:

establishing a pricing tier for each request generating device, a maximum rate of request servicing, and an expected time for serving a request at the maximum rate of request servicing;

for each request generating device with a premium pricing tier, maintaining an instantaneous rate of request servicing by the request servicing device;

following servicing of each request from a request generating device by the request servicing device,

determining a time elapsed during servicing of the request;

when the time elapsed during servicing of the request is less than the expected time for serving a request established for the request generating device,

calculating a remaining time equal to the difference between expected time for serving a request established for the request generating device and the time elapsed during servicing of the request; and

waiting for a length of time based on the calculated remaining time prior to servicing another request for the request generating device.

2. (original) The method of claim 1 wherein the length of time based on the calculated remaining time is determined to be:

the calculated remaining time for a request generating device for which the established pricing tier is a basic pricing tier;

the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is equal to the maximum rate of request servicing established for the request generating device;

greater than the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is greater than the maximum rate of request servicing established for the request generating device; and

less than the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is less than the maximum rate of request servicing established for the request generating device.

3. (original) The method of claim 2 wherein the request generating device is a computer.

4. (original) The method of claim 2 wherein the request servicing device is an electronic data storage device.

5. (original) The method of claim 4 wherein the electronic data storage device is a disk array.

6. (currently amended) The method of claim ~~6~~ 1 wherein the maximum rate of request servicing is established via specification of a maximum rate of request servicing by the request generating device.

7. (original) The method of claim 1 wherein the maximum rate of request servicing is established by partitioning the capacity of the request servicing device among the request generating devices in order to provide, when possible, each request generating device with a maximum rate of request servicing specified by the request generating device, and otherwise to provide each request generating device with a maximum rate of request servicing proportional to a maximum rate of request servicing specified by the request generating device.

8. (original) The method of claim 1 wherein the request servicing device may dynamically alter the maximum rate of request servicing provided to one or more request generating devices in accordance with a rate at which the request servicing device receives requests and according to the request servicing capacity of the request serving device.

9. (original) The method of claim 1 wherein maintaining an instantaneous rate of request servicing by the request servicing device further comprises:

initially setting the instantaneous rate of request servicing for a request generating device to one request divided by the expected time for serving a request at the maximum rate of request servicing established for the request generating device;

increasing the instantaneous rate of request servicing for the request generating device by one following servicing of a request generated by the request generating device; and

decreasing the instantaneous rate of request servicing for the request generating device by one at regular intervals of time.

10. (original) A request servicing device that fairly services electronic requests received by the request servicing device from request generating devices interconnected with the request receiving device, the request servicing device including:

a memory that contains an established maximum rate of request servicing, an expected time for serving a request at the maximum rate of request servicing, and a pricing tier for each request generating device and that contains, for each request generating device with a premium pricing tier, an instantaneous rate of request servicing by the request servicing device; and

control functionality that services electronic requests received from the request generating devices and that, following servicing of each request from a request generating device by the request servicing device, determines a time elapsed during servicing of the request so that, when the time elapsed during servicing of the request is less than the expected time for serving a request established for the request generating device, the control functionality calculates a remaining time equal to the difference between expected time for serving a request established for the request generating device and the time elapsed during servicing of the request and waits for a length of time based on the calculated remaining time prior to servicing another request for the request generating device.

11. (original) The request servicing device of claim 1 wherein the length of time based on the calculated remaining time is determined by the request servicing device to be:

the calculated remaining time for a request generating device for which the established pricing tier is a basic pricing tier;

the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is equal to the maximum rate of request servicing established for the request generating device;

greater than the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is greater than the maximum rate of request servicing established for the request generating device; and

less than the calculated remaining time for a request generating device for which the established pricing tier is a premium pricing tier and the instantaneous rate of request servicing is less than the maximum rate of request servicing established for the request generating device.

12. (original) The request servicing device of claim 10 wherein the request generating device is a computer.

13. (original) The request servicing device of claim 10 wherein the request servicing device is an electronic data storage device.

14. (original) The request servicing device of claim 10 wherein the electronic data storage device is a disk array.

15. (original) The request servicing device of claim 10 wherein the maximum rate of request servicing is established via specification of a maximum rate of request servicing by the request generating device.

16. (original) The request servicing device of claim 10 wherein the maximum rate of request servicing is established by partitioning the capacity of the request servicing device among the request generating devices in order to provide, when possible, each request generating device with a maximum rate of request servicing specified by the request generating device, and otherwise to provide each request generating device with a maximum rate of request servicing proportional to a maximum rate of request servicing specified by the request generating device.

17. (original) The request servicing device of claim 10 wherein the request servicing device may dynamically alter the maximum rate of request servicing provided to one or more request generating devices in accordance with a rate at which the request servicing device receives requests and according to the request servicing capacity of the request serving device.

18. (original) The request servicing device of claim 10 wherein maintaining an instantaneous rate of request servicing by the request servicing device further comprises:

initially setting the instantaneous rate of request servicing for a request generating device to one request divided by the expected time for serving a request at the maximum rate of request servicing established for the request generating device;

increasing the instantaneous rate of request servicing for the request generating device by one following servicing of a request generated by the request generating device; and

decreasing the instantaneous rate of request servicing for the request generating device by one at regular intervals of time.

19. (original) The request servicing device of claim 18 wherein separate execution threads in a firmware or software implementation of control functionality within the request servicing device increase the instantaneous rate of request servicing and decrease the instantaneous rate of request servicing.

20. (original) The request servicing device of claim 10 wherein the length of time based on the calculated remaining time determined by the request servicing device to be greater than the calculated remaining time for a request generating device is further determined to be a ratio multiplied by the calculated remaining time, the ratio calculated by dividing the instantaneous rate of request servicing by the expected time for serving a request, and wherein the length of time based on the calculated remaining time determined by the request servicing device to be less than the calculated remaining time for a request generating device is further determined to be a ratio multiplied by the calculated remaining time, the ratio calculated by dividing one by the difference between the expected time for serving a request and the instantaneous rate of request servicing.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.